

AI/ML for Data and Analytics

Ramesh Natarajan

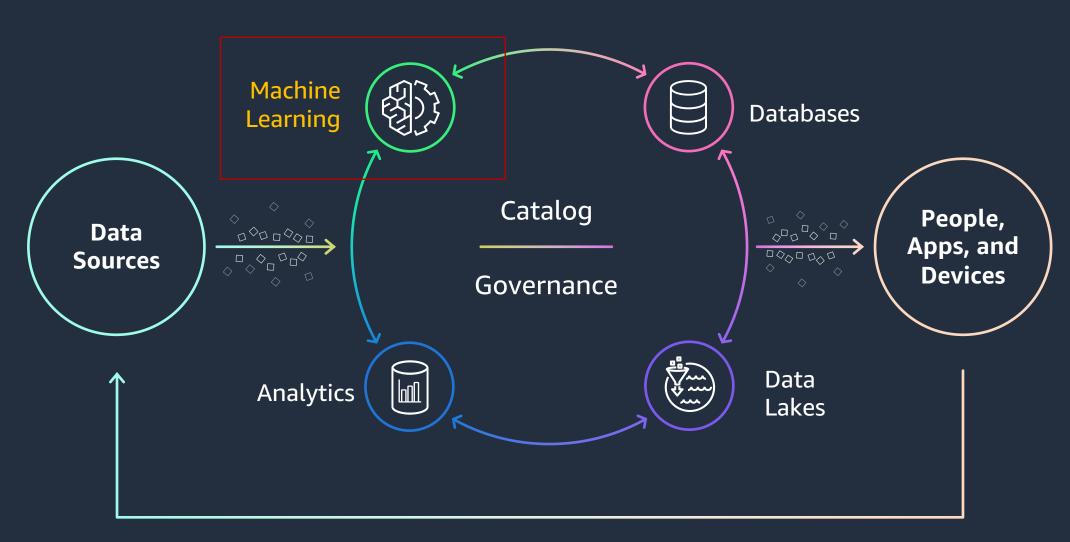
Solutions Architect Manager natarram@amazon.com

Agenda

- AI/ML in modern data strategy
- AI/ML use cases for public sector
- AWS AI/ML Stack
 - Generative AI use cases for public sector
 - Amazon Bedrock, Amazon Q in QuickSight



AI/ML in modern data strategy





AI/ML in end-to-end modern data foundation

Data sources

OLTP SOURCES

STREAMING DATA

THIRD-PARTY DATA



Amazon Aurora

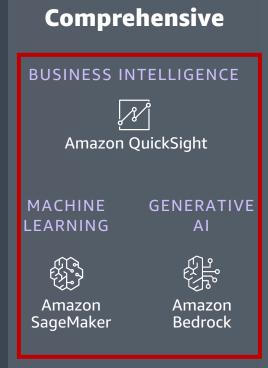
DATABASES

Amazon Redsh

Amazon DynamoDB

Amazon Kinesis & Amazon MSK







Catalog and govern

AWS Lake Formation, Amazon DataZone



AI, ML, Deep learning?



Artificial intelligence (AI)

Any technique that allows computers to mimic human intelligence using logic, if-then statements, and machine learning



Machine learning (ML)

A subset of AI that uses machines to search for patterns in data to build logic models automatically



Deep learning (DL)

A subset of ML composed of deeply multi-layered neural networks that perform tasks like speech and image recognition



Generative Al

Powered by large models that are pretrained on vast corpuses of data and commonly referred to as foundation models (FMs)

Challenges we are hearing from public sector customers



Demand for government services is rising while resources and capacity to deliver them **aren't keeping pace**



Citizens increasingly expect government to **provide modern digital experiences** for conducting online transactions



Aging infrastructure for data capture, storage, and management **creates friction** for leveraging data for analytics and machine learning

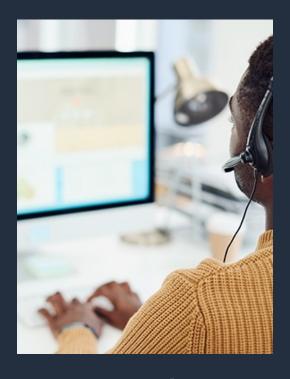


Complex security, privacy, and compliance requirements create barriers to change and block adoption of many SaaS solutions



Risk averse culture and institutional inertia slow innovation

Top AI/ML use cases for public sector



Speech and language

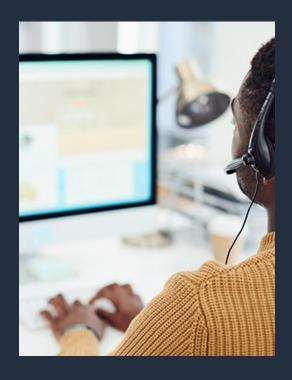


Intelligent document processing



Predictions and insights

Top AI/ML use cases for public sector



Speech and language



Intelligent document processing



Predictions and insights

AI/ML-Enabled Citizen (or Student) engagement

Engage citizens (or students) and drive improvements in customer satisfaction

- Improve contact center agent effectiveness with real-time translation and decision support using Amazon Connect and Contact Center Intelligence
- Analyze call and text interactions with citizens to spot issues and trends and drive improvement
- Improve self service



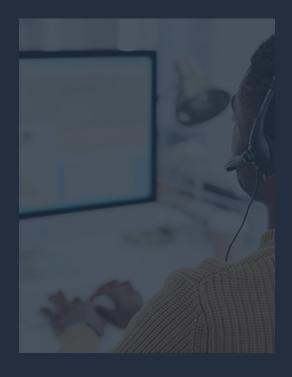
Using Al to improve agent efficiency

"During peak hours, previously you're 45-50 minutes on hold, and now that's has been reduced to about three and a half minutes. One of the other benefits we've gotten from Amazon Connect is sentiment analysis. On a call, we get real-time feedback on whether or not the customer was happy, frustrated, or angry..."

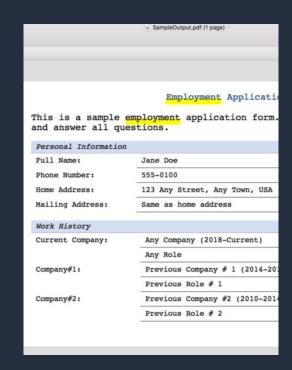
—Benny Chacko, Deputy General – LA County Internal Services Department



Top AI/ML use cases for public sector



Speech and language



Intelligent document processing

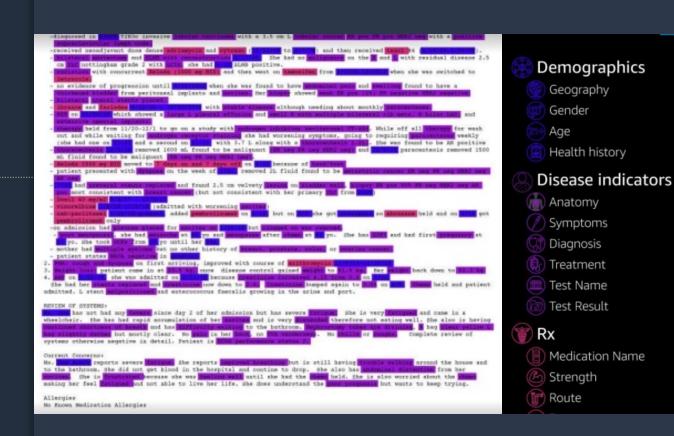


Predictions and insights

Extract insights from unstructured content

Extract insights from unstructured documents and forms, like images, PDFs, and audio

- Analyze text with natural language processing (NLP) to identify topics, extract entities, understand sentiment, and classify documents with Amazon Textract, Amazon Rekognition, and Amazon Comprehend
- Translate content at scale with Amazon Translate



King County Assessor's Office



CHALLENGE:

Reduce data entry, eliminate data errors, and improve data time lines.

SOLUTION:

Intelligent documents processing for documents and electronic files, streamlining & unlock data and information from paper documents and electronic files

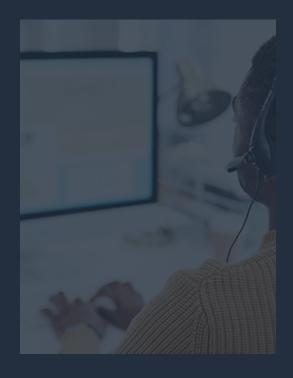
RESULTS:

King County employees will focus on higher value, more satisfying work, and ultimately help the County realize its vision for connected communities, connected data, and connected government."



This Photo by Unknown Author is licensed under CC BY-SA

Top AI/ML use cases for public sector



Speech and language



Intelligent document processing

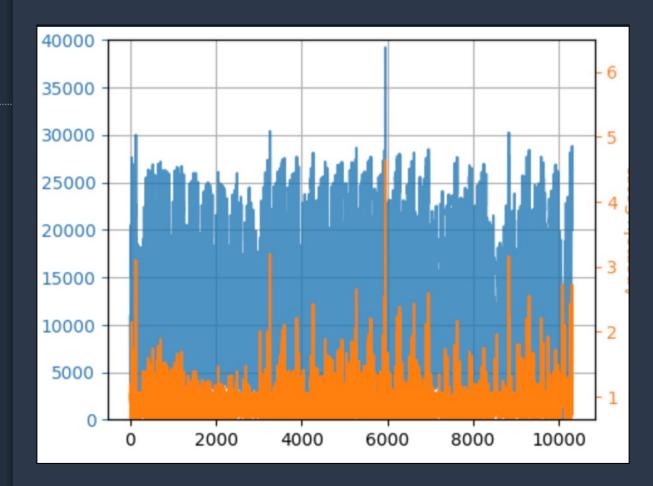


Predictions and insights

Fraud detection and prevention

Detect and prevent fraud, waste, and abuse

- Enhance accuracy and speed to help detect and prevent waste fraud and abuse
- Managed service approach with prebuilt
 ML models for fraud detection
- Supervised and unsupervised models for developing highly targeted models to utilize customer data as part of fraud prevention efforts

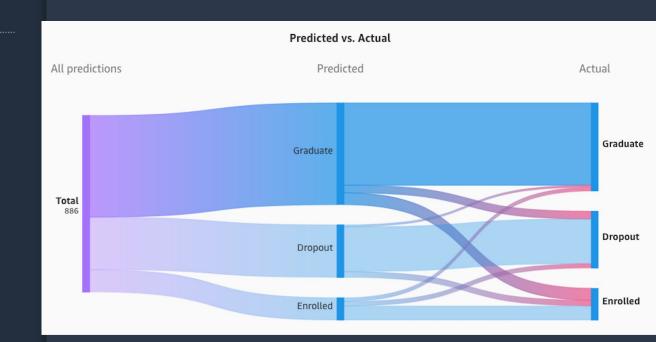




Machine learning to forecast trends and support decisions

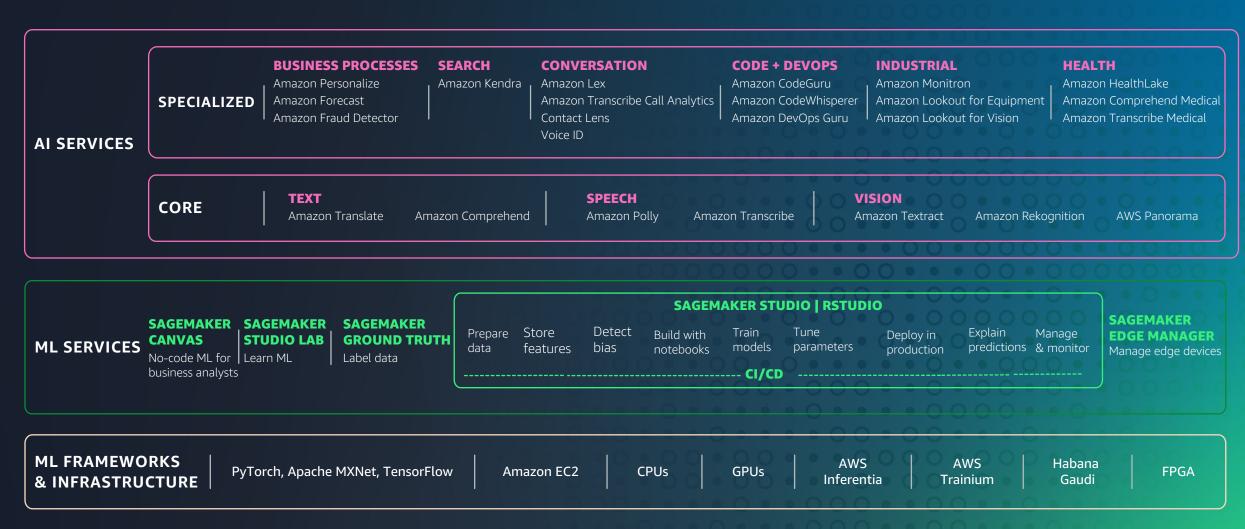
Prediction and forecast with machine learning

- Predicting service demand or program activities
- Allocating resources to optimize impact and outcomes for citizens
- Financial planning and revenue / cost forecasts
- Predicting the academic outcomes of higher-education students



The AWS AI/ML stack

Broadest and most complete set of machine learning capabilities

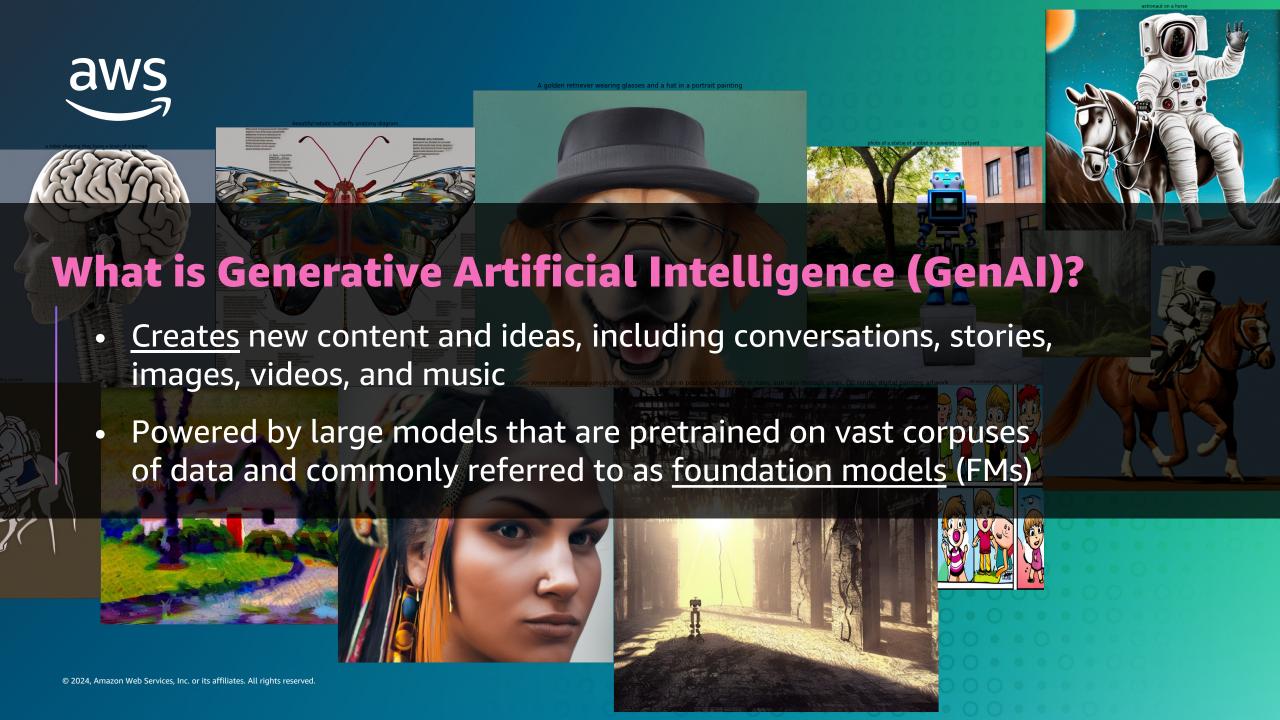


Generative AI is top of mind for most all of our customers

Key drivers: Model size | Transformer technology | Ease of access

Expectations: Massive productivity gains | Innovations @ scale





Customize FMs Emerging generative AI data patterns

1



Retrieval Augmented Generation (RAG)

Guide pre-trained models with private, domain-specific contextual data

For example, virtual agents with limited domain-specific requirements

2



Fine-tune a pre-trained model

Further training of pre-trained models with domain-specific data

For example, domain-intensive knowledge agents

3



Train your own model

Training of models on domain-specific data

For example, deep domain-specific applications trained on specific data



Generative Al Stack

APPLICATIONS THAT LEVERAGE LLMs AND OTHER





Amazon Q in Amazon Q in Amazon Q in Amazon Conn



Amazon Connect



Amazon CodeWhisperer

TOOLS TO BUILD WITH LLMs AND OTHER FMs



Amazon Bedrock



Guardrails | Agents | Customization Capabilities

INFRASTRUCTURE FOR FM TRAINING A









GPUs Trainium La Inferentia SageMaker







UltraClusters EFA EC2 Capacity Blocks Nitro Neuron





Common Generative Al use cases



Text generation



Q&A



Text summarization



Text extraction



Paraphrase rephrase



Search



Code generation



Image generation



Image classification



Audio generation



Video generation



Generative AI public sector application examples



Constituent Communications

Citizen engagement and feedback, transparency



Personalization

Create personalized learning pathways for students



Public Health

Personalized care, population health assessments



Productivity

Transform information into sample test questions and accelerate grading



Public Safety

Public safety and crime prevention, Emergency response and disaster management



Virtual assistants

Enhance the student experience with humanlike responses



Transportation

Traffic optimization, autonomous vehicle control, personalized transportation experiences

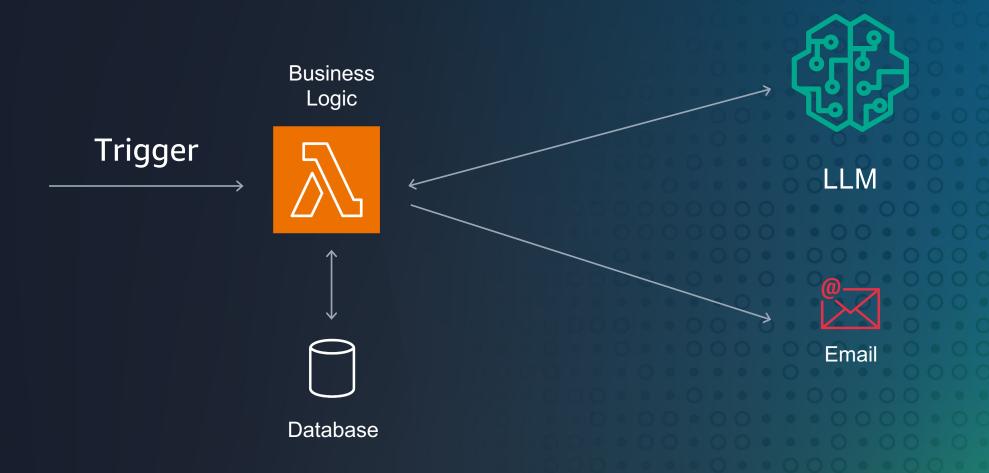


Content generation and summarization

Generate course materials, images, videos, and music. Summarize research documents, lecture transcripts, and class notes



Use Case: Program Enrollment Notifications (Realtime)





Use Case: Customized Insights (Batch)

Business Logic Periodically LLM Scheduled Job Alert Data lake / warehouse





Amazon Bedrock

The easiest way to build and scale generative AI applications with foundation models (FMs)

Choice of leading FMs through a single API

Model customization

Retrieval Augmented Generation (RAG)

Agents that execute multistep tasks

Security, privacy, and safety



Amazon Bedrock Data Protection Helps keep your data secure and private

None of the customer's data is used to train the underlying model

All data is encrypted in transit and at rest; data used for customization is securely transferred through customer's VPC

Data remains in the Region where the API is processed

Support for GDPR, SOC, ISO, CSA compliance, and HIPAA eligibility





Guardrails for Amazon Bedrock

Implement safeguards customized to your application requirements and responsible AI policies

Apply guardrails to multiple foundation models and Agents for Amazon Bedrock

Configure harmful content filtering based on your responsible AI policies

Define and disallow denied topics with short natural language descriptions

Redact or block sensitive information such as PIIs, and custom Regex.

Amazon QuickSight

- Create beautiful, interactive dashboards
- Add rich interactivity like filters, drill downs, zooming, and more
- Blazing fast navigation
- Accessible on any device
- Data Refresh
- Publish to everyone with a click





Amazon Q in QuickSight capabilities

Accelerate the path to data-driven decisions for ALL USERS

AI-powered dashboard authoring experience enables business analysts to build faster

AI-assisted data storytelling

enables business users to discover and share findings

Al answers to questions of data on demand simplifies extracting insights for business users

Extend custom apps with AI

enables developers to quickly add advanced capabilities

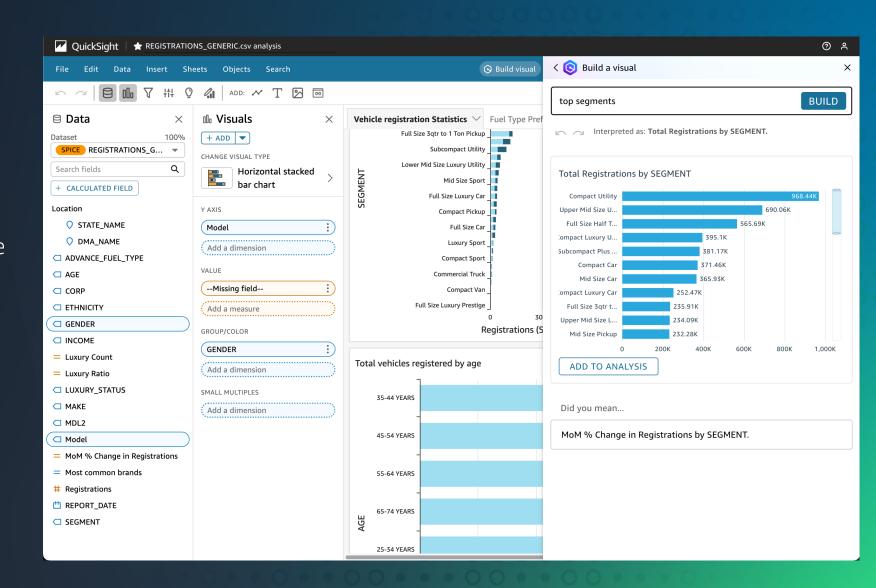
Powered by Amazon Bedrock



Al-powered dashboard authoring experience

Amazon Q in QuickSight

Use everyday language to generate and fine-tune visuals in seconds

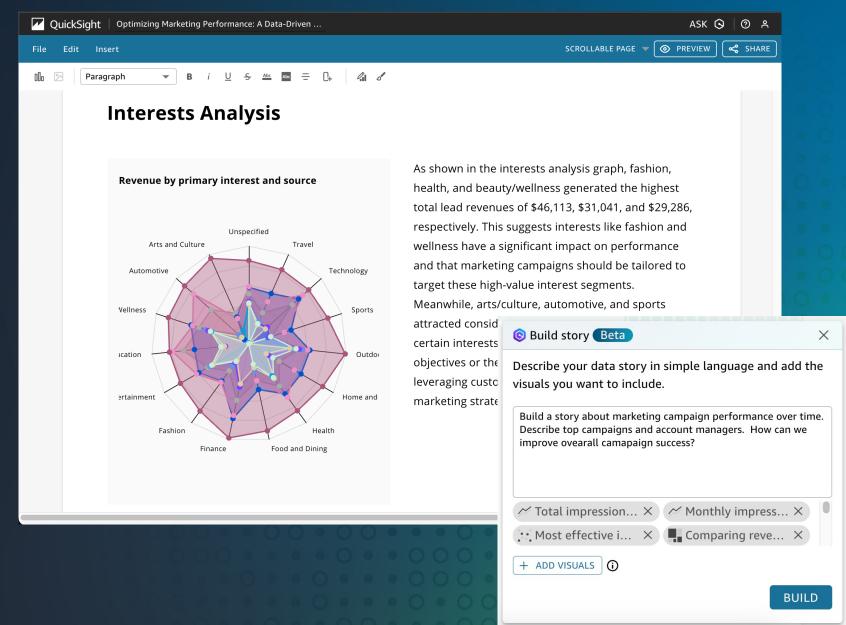




Al-assisted data storytelling

Amazon Q in QuickSight

Automatically generate data stories using natural language prompts



Call to action - AI/ML for Data and Analytics in Public Sector

- How can we support you in ensuring accuracy and authority of model outputs?
- How can we use "guardrails" to minimize inappropriate content?
- How can we maximize the public investment and minimize cost?
- What are your latency requirements for specific use cases?
- What are the potential regulatory, data privacy, & security considerations that may dictate model architecture.
- How can we ensure proprietary, copyrighted and IP concerns are addressed?



Thank you!

Ramesh Natarajan

natarram@amazon.com

Please Provide Your Feedback



Step 1: Select "Data and Analytics Track"

Step 2: Select "AI/ML for Data & Analytics"

Learning Day Content

https://sanfrancisco2024.awslearningday.com/



